

An AI Approach to Ground Station Autonomy for Deep Space Communications

Forest Fisher, Tara Estlin, Darren Mutz, and Steve Chien

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, M/S 126-347, Pasadena, CA 91109-8099
phone: +1 818 393-5368, fax: +1 818 393-5244, email: {firstname.lastname}@jpl.nasa.gov

ABSTRACT

This paper describes components of a system for an autonomous deep space tracking station. The system enables fully automated routine operations encompassing scheduling and resource allocation, antenna and receiver predict generation, track procedure generation from service requests, and closed loop control and error recovery for the station subsystems. This system has been validated by the construction of a prototype Deep Space Terminal (DS-T) tracking station, which has performed a series of demonstrations of autonomous ground station control for downlink services with NASA's Mars Global Surveyor (MGS).

INTRODUCTION

The Deep Space Network (DSN) [8] was established in 1958 and has since evolved into the largest and most sensitive scientific telecommunications and radio navigation network in the world. The purpose of the DSN is to support unmanned interplanetary spacecraft missions and to support radio and radar astronomy observations taken in the exploration of space. The DSN currently consists of three deep-space communications facilities placed approximately 120 degrees apart around the world: at Goldstone, in California's Mojave Desert; near Madrid, Spain; and near Canberra, Australia. This strategic placement permits constant observation of spacecraft as the Earth rotates, and helps to make the DSN the largest and most sensitive scientific telecommunications system in the world. Each DSN complex operates a set of deep space stations consisting of 70-meter, 34-meter, and 26-meter antennas. The function of the DSN is to receive telemetry signals from spacecraft, transmit commands that control spacecraft operating modes, generate the radio navigation data used to locate and guide a spacecraft to its destination, and acquire flight radio science, radio and radar astronomy, very long baseline interferometry (VLBI), and geodynamics measurements.

From its inception the DSN has been driven by the need to create increasingly more sensitive telecommunications devices and better techniques for navigation. The operation of the DSN communications complexes requires a high level of manual interaction with the devices in the communications link with the spacecraft. In more recent times NASA has added some new drivers to the development of the DSN:

1. reduce the cost of operating the DSN,
2. improve the operability, reliability, and maintainability of the DSN, and
3. prepare for a new era of space exploration with the New Millennium program: support small, intelligent spacecraft requiring very few mission operations personnel [14].

In the past, the process of operating such stations has been labor and knowledge intensive. Recently, efforts have been made to reduce the cost of operations. One such effort has been in the area of antenna station automation. Many approaches have been applied to automation control/commanding of different types of antenna systems. In the AI group at JPL, we have worked on automating the scheduling of communications antennas and the generation of antenna command sequences. The scheduling of communications antennas consist of allocating an oversubscribed resource, the antenna, to a flight project in order to provide communication services, while antenna command sequences set up and perform a particular communications link with a spacecraft [2]. These sequences can be run as control scripts to operate the station and all of its relevant subsystems [14]. This work was demonstrated as a component of the Deep Space Terminal (DS-T) during a series of demonstrations from April to September of 1998. Through the use of these technologies a high level goal-oriented interface is provided to the system. This interface enables users to specify *what* they want done and does not require that they specify or even know *how* it should be done.

The rest of this paper is organized in the following manner. We first offer a brief background on how the DSN operates. Next we provide an introductory explanation of the DS-T functionality. From here we discuss two of the underlying technologies providing much of the DS-T's autonomy: automated scheduling and automated planning. We then conclude the paper with results from our demonstrations and talk about future work.

How the DSN Works

The DSN track process occurs daily for dozens of different NASA spacecraft and projects, which use the DSN to capture spacecraft data. Though the process of sending signals from a spacecraft to Earth is conceptually simple, in reality there are many earthside challenges that must be addressed before a spacecraft's signal is acquired and successfully transformed into useful information. In the remainder of this section, we outline some of the steps involved in providing tracking services and in particular discuss the problem of track plan generation.

The first step in performing a DSN track is called *network preparation*. Here, a project sends a request for the DSN to track a spacecraft involving specific tracking services (e.g., downlink, uplink). The DSN responds to the request by attempting to schedule the necessary resources (i.e. an antenna and other shared equipment) needed for the track. Once an equipment schedule and other necessary information has been determined, the next step is the *data capture process*, which is performed by operations personnel at the deep space station. During this process, operators determine the correct steps to perform the following tasks: configure the equipment for the track, perform the actual establishment of the communications link, and then perform the actual track by issuing control commands to the various subsystems comprising the link.

Throughout the track the operators continually monitor the status of the link and handle exceptions (e.g., the receiver breaks lock with the spacecraft) as they occur. All of these actions are currently performed by human operators, who manually issue hundreds of commands via a computer keyboard to the link subsystems.

This paper discusses the application of two AI systems for automated antenna operations. These systems are an AI scheduling system for allocating communications resources, and an AI planning system to generate deep space communication antenna control scripts. These two

components are intended to dramatically reduce the need for many manual steps.

Deep Space Terminal

The components discussed in this paper were demonstrated as part of the Deep Space Terminal (DS-T), a prototype 34-meter deep space communications station developed [9][10][11] as a technology demonstration of fully autonomous *lights-out* operations. In the DS-T concept, a global DSN schedule is disseminated to a set of autonomous DS-T stations, where each DS-T station operates autonomously, performing tracks in a largely independent fashion. When requested to perform a track, the DS-T station performs a number of tasks (at appropriate times) required to execute the track. First, the DS-T station uses appropriate spacecraft navigation ephemeris and predict generation software in order to produce necessary antenna and receiver predict information required to perform the track. Next, the DS-T station executes the pre-calibration process, in which the antenna and appropriate subsystems (e.g., receiver, exciter, telemetry processor, etc.) are configured in anticipation of the track. During the actual track, the signal from the spacecraft must be acquired and the antenna and subsystems must be commanded to retain the signal, adjust for changes in the signal (such as changes in bit rate or modulation index as transmitted by the spacecraft), and perform error recovery. Finally, at the completion of the track, the station must be returned to an appropriate standby state in preparation for the next track. All of these activities require significant automation and robust execution including closed loop control, retries and contingency handling.

In order to provide this autonomous operation capability, the DS-T station employs tightly coupled state of the art hardware and software. At the core of the autonomy are two areas of artificial intelligence (AI) technology, AI scheduling and AI planning. We will offer a brief example of each and a brief context for how they apply to the DS-T.

The original goal of the DS-T task was to build an autonomous control system for a deep space communications station. This system had to meet the following criteria: schedule driven with a high level service request interface; an automated scheduling component for initial scheduling and rescheduling; provide script guided control; ability to generate predicts or use provided predicts; automatically configure pre-track; utilization of COTS (Commercial Off The Shelf) components wherever feasible; operations based on

defined but expandable set of services; autonomous error recovery for a defined class of problems; post pass data



Figure 1: 34m BWG Antennas at Goldstone

delivery; and treat ground terminal as a network computer with an RF peripheral.

One of the most important points was the idea of a ground station looking just like a network computer to a user, operator, or mission. This is best demonstrated by an operational scenario. To provide service a user need only login to the DS-T work-station and submit a service request to the scheduling system, or FTP a schedule and service request to a particular file system location. From either of these inputs DS-T would detect the existence of a track/service schedule, proceed to schedule station specific tasks, configure the station to provide the service, and finally when the time comes, the track would begin without further user interaction.

As mentioned above, the station reacts to a service request derived schedule generated by an automated scheduling system. It is through the reaction to this schedule that the dynamic track-specific control scripts are generated. Autonomous operations of the station takes place through the execution of these control scripts.

In Figure 1, we show a picture of the three 34-meter Beam Wave Guide antennas at Goldstone, CA. In the foreground is DSS-26, which was the station selected for prototyping the DS-T.

In April 1998, the DS-T prototype first demonstrated automated downlink capability of single isolated tracks for the Mars Global Surveyor (MGS) spacecraft. Between April and September 1998, many multi-day demonstrations took place including a six day unattended demonstration. During these demonstrations, a service request for downlink services, a track sequence of events, and spacecraft ephemeris data were used to automatically downlink data from the MGS spacecraft.

Scheduling for DS-T

When the decision is made to fly a mission, a very knowledge-intensive process begins that will ensure the necessary DSN antenna coverage. First, a forecast is made of the DSN resources that the spacecraft will require. In the Resource Allocation Process (RAP), the types of services, frequency, and duration of the required tracks are determined as well as high-level re-source requirements (e.g., antenna). While the exact timing of the tracks is not known, a set of automated forecasting tools are used to estimate network load and to assist in ensuring that adequate network resources will be available. One part of the network architecture is a unified tool suite that has been developed called TMOD Integrated Ground Resource Allocation System (TIGRAS), which uses operations research and probabilistic reasoning techniques to allow forecasting and capacity planning for DSN resources [1].

As the time of the actual tracks approaches, this estimate of resource loading is converted to an actual schedule, which becomes more concrete as time progresses. In this process, specific project service requests and priorities are matched up with available resources in order to meet communications needs for earth-orbiting and deep space spacecraft. This scheduling process involves considerations of thousands of possible tracks, tens of projects, tens of antenna resources and considerations of hundreds of subsystem configurations. In addition to adding the detail of antenna subsystem allocation, the initial schedule undergoes continual modification due to changing project needs, equipment availability, and weather considerations. Responding to changing context and minimizing disruption while rescheduling is a key issue.

At the high level of resource allocation, schedule execution does not involve execution monitoring. However, rescheduling is often necessary due to: equipment outages, last minute track requests, last minute changes to scheduled tracks, and changing atmospheric conditions. Rescheduling can occur in two ways: (1) it can be initiated top-down due to a change to a previously scheduled track or the addition of another request; and (2) it can occur bottom-up in that equipment outages can occur or tracks can fail necessitating rescheduling. From the standpoint of the scheduler the important feature is the degree of change required to make the schedule consistent.

Because of the size and complexity of the rescheduling task, manual scheduling is prohibitively expensive.

Automation of these scheduling functions is projected to save millions of dollars per year in DSN operations costs. Based on these motivating factors, the Demand Access Network Scheduler (DANS), which was designed to deal with the complex subsystem and priority schemes required to schedule the 34 and 70 meter antennas, was used as one of the scheduling components of the DS-T.

DANS: Automated Scheduling

The Demand Access Network Scheduler (DANS) [3] system, designed to deal with the complex subsystem and priority schemes required to schedule the larger 34 and 70 meter antennas, uses the forecasted antenna schedule produced by the RAP process and supports rescheduling as required by changing tracking requirements and equipment availability. The main inputs to DANS are the current schedule and a set of new tracking requests or changes to current tracks and/or equipment that must be handled in the final schedule. A tracking request usually specifies information such as the spacecraft or project name (e.g., DS1, Voyager), the type of antenna requested (e.g., 70M, 34M), the number of individual tracks requested (e.g., 4 tracks per week), the start time and end time for each track, priority of each track, etc.

DANS uses priority-driven, best-first, constraint-based search and iterative optimization techniques to perform priority-based rescheduling in response to changing network demand. In this approach, DANS first considers the antenna allocation process, since antennas are the central focus of resource contention. After establishing a range of antenna options, DANS then considers allocation of the 5-13 subsystems per track (out of the tens of shared subsystems at each antenna complex) used by each track. DANS uses constraint-driven, branch and bound, best-first search to efficiently consider the large set of possible subsystems schedules. The DANS objective is to satisfy as many activity requests as possible while maintaining a conflict-free status (i.e. no hard constraints violated) with minimal disruption to the existing schedule.

The DSN scheduling problem is complicated by three factors: (1) context-dependent priority; (2) subsystem allocation; and (3) the possibility of reducing the length of the tracks. DSN track priorities are context dependent in that they are often contingent on the amount of tracking the project has received so far in the week. For example, a project might have priority 3 to get 5 tracks, priority 4 to get 7 tracks and priority 6 to get 9 tracks (where lower priority numbers represent more important tracks). This reflects that 5 tracks are necessary to maintain spacecraft health and get critical science data to

ground stations; 7 tracks will allow a nominal amount of science data to be downlinked; and 9 tracks will allow for downlinking of all science data (e.g., beyond this level additional tracks have little utility). An important point is that specific tracks are not labeled with these priorities (e.g., the project is allowed to submit 5 tracks at priority 3, 2 at priority 4 and so on). Rather, when considering adding, deleting, or moving tracks the scheduler must consider the overall priority of the project in the current allocation context.

Finally, the DSN scheduling problem is complicated by the fact that the track duration can be relaxed. For example, a project may request a 3 hour track but specify a minimum track time of 2 hours. When evaluating potential resource conflicts the scheduler must consider the option of shortening tracks to remove resource conflicts. DANS also uses a linear weighting scheme in conjunction with a modified SIMPLEX algorithm to trim tracks in accordance with prioritizations.

Once generated, a schedule is first used at a network wide level designating what resources (primarily the antennas) shall be used to provide what services (primarily communications tracks). In the DS-T architecture the schedule is then disseminated to each DS-T station to designate when and what type of service is to be performed by that station. From this high level description of the service, each DS-T station proceeds to schedule station specific activities in order to provide the desired services. This secondary station specific scheduling component utilized a simple macro expansion scheduling algorithm and controls the execution of the schedule. These activities consist of track script generation and execution of the track script for each track.

Planning for DS-T

Once a DS-T station has been allocated to provide the communications service for a particular mission (i.e. has been scheduled) and the station specific activities have been scheduled, the DS-T script generator is invoked just prior to the track beginning in order to generate the antenna control script. The final result is the set of antenna commands necessary to setup and perform the request track (communication service).

The DS-T script generator (SG) is where the majority of the control autonomy is provided. The SG uses Artificial Intelligence planning techniques to perform a complex software module reconfiguration process [5]. This process consists of piecing together numerous highly interdependent smaller control scripts in order to produce

a single script to control the operations of the DS-T station.

The core engine used in the SG is the Automated Scheduling and Planning ENvironment (ASPEN) [13]. The ASPEN system is a reusable, configurable, generic planning/scheduling application framework that can be tailored to specific domains to create conflict-free plans or schedules. It has a number of useful features including an expressive modeling language, a constraint management system for representing and maintaining antenna operability and/or resource constraints, a temporal reasoning system and a graphical interface for visualizing plans and states. ASPEN has been adapted to input antenna-tracking goals and automatically produce the required command sequence necessary to create the requested link [12].

The control script produced by the SG:

- sets up the track by configuring the station during pre-track;
- provides the track service requested by commanding the antenna and sub-systems to acquire and maintain lock on the signal throughout mode changes; and
- cleans up and shuts down the station at the completion of the track.

TRACK PLAN GENERATION: THE PROBLEM

Generating an antenna track plan involves taking a general service request (such as telemetry - the downlink of data from a spacecraft), an antenna knowledge-base (which provides the information on the requirements of antenna operation actions), and other project specific information (such as the spacecraft sequence of events), and then generating a partially-ordered sequence of commands. This command sequence will properly configure a communications link that enables the appropriate interaction with the spacecraft. To automate this task, the ASPEN planning and scheduling system has been applied to generate antenna operation procedures on demand.

ASPEN has been adapted to use high-level antenna track information to determine the appropriate steps, parameters on these steps and ordering constraints on these steps that will achieve the input track goals. In generating the antenna track plan, the planner uses information from several sources. In Figure 2 we show the inputs and output of the DS-T script Generator.

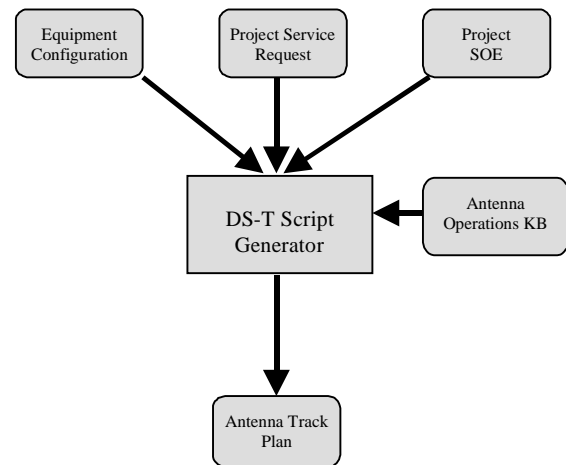


Figure 2: DS-T Script Generator Inputs and Outputs

Equipment Configuration - This configuration details the types of equipment available and includes items such as the antenna, antenna controller, the receiver, etc.

Project Service Request - The service request specifies the DSN services (e.g., downlink, uplink) requested by the project and corresponds to the goals or purpose of the track.

Project SOE - The project sequence of events (SOE) details spacecraft events occurring during the track - including the timing of the beginning and ending of the track and spacecraft data transmission bit rate changes, modulation index changes, and carrier and subcarrier frequency changes.

Antenna Operations KB - The Antenna Operations Knowledge Base (KB) stores information on available antenna operations actions/commands. This KB dictates how actions can be combined to provide essential communication services. Specifically, this includes information such as action preconditions, postconditions, and command directives and also includes any other relevant information such as resource and state descriptions.

Antenna Track Plan - The Antenna Track Plan is the output of the ASPEN/DS-T Script Generator. The track plan is the dynamically produced command control script. When executed these scripts issue all of the necessary subsystem command directives to configure, control, and perform the communications track.

Through the use of the ASPEN/DS-T Script Generator these high level inputs provide the goal-oriented interface

enabling the system to be used by specifying *what* is to be done instead of *how* it should be done.

DS-T Demonstrations

The Deep Space Terminal (DS-T) [10][11] concept was validated through a number of demonstrations. These began with the automation of partial tracks in April 1998, continued with 1-day unattended operations in May, and concluded with a 6-day autonomous “lights-out” demonstration in September 1998. Throughout these demonstrations ASPEN was used to automatically generate the necessary command sequences for a series of Mars Global Surveyor (MGS) downlink tracks using the equipment configuration at Deep Space Station 26 (DSS26), a 34-meter antenna located in Goldstone, CA. These command sequences were produced and executed in a fully autonomous fashion with no human intervention. During the September demonstration, DS-T performed all Mars Global Surveyor coverage scheduled for the Goldstone antenna complex. This corresponded to roughly 13 hours of continuous track coverage per day.

While the overall DS-T effort consisted of a large team and a project duration of approximately 1.5 years, the DS-T automation team consisted of three team members. Of this team’s work, approximately one work year was spent on the script generation effort. This effort primarily consisted of knowledge acquisition and model development, while a small effort was made in the integration of the script generator. A key factor in the quick development was the ability to adapt a general purpose planning and scheduling system. As the domain of ground communication-station commanding shared many similarities to spacecraft commanding, ASPEN seemed like a logical choice. This was confirmed by the ease of knowledge base development and integration. Spacecraft commanding also consists of generating a sequence of commands, however it is predominately a resource-scheduling problem, whereas ground-station commanding is predominately a sequencing problem.

Results

In order to provide qualitative results, we present statistical data from September 16, 1998, a representative day during our 6-day autonomous unattended demonstration, during which we collected above 90% of the transmitted frames. This performance is on par with the operator-controlled stations, however required no support personnel (i.e. reduced operations cost).




Figure 3: September 16, 1998 MGS Track

In Figure 3, the graph represents when MGS was in view of the ground stations at each of the three complexes (Madrid, Goldstone, and Canberra). DS-T, which is located at Goldstone, tracked MGS through the five track segments indicated in Figure 3.

Before continuing with the analysis of the results, let us explain the different modes indicated in Figure 3 for each of the different track segments. When a spacecraft is downlinking data it is said to be in 1way mode. When an uplink and a downlink are taking place simultaneously the spacecraft is said to be in 2way mode. If a station is communicating in 2way mode with a spacecraft, and another station is listening in on the downlink of the spacecraft, the second station is said to be in 3way with the 2way station. Because DS-T is not equipped for uplink, DS-T operates in either 1way or 3way mode. In this example, during segment 4 dss25 (deep space station) was in 2way and DS-T was in 3way with 25 (3way/25).

Track segment 2, which is labeled LOS, indicates that there was a scheduled loss of signal (LOS) so during this segment no frames were collected. During each of the other respective track segment DS-T collected 75%, 91%, 96%, 90%, 23% of the broadcasted frames. As shown by the graph, during segment 1 and 6 the elevation of the dish is low in the sky. Under these circumstances there is considerably more atmospheric interference which explains the lower percent of frame collection. On the other hand, if you look at segment 4, where there is a long segment with the spacecraft high in the sky, the data collection is quite high. In segment 3 and 5 the values are a little lower due to the shortness of the segments.

This is explained by the fact that some data is lost during a change in mode, as in the transition from LOS to 1way and 3way/25 to 1way.

As a component of the DS-T, the DANS automated scheduler enabled us to demonstrate how a network of DS-T like terminals would perform in a schedule driven environment. It is partially through this functionality of starting from a high-level service request and producing a resource allocation schedule that the DS-T concept is able to provide communications service through a high level interface. In conjunction with the scheduling system the DS-T Script Generator performed flawlessly, producing dynamically instantiated control scripts based on the desired service goals for the communications pass as specified in the service request. The use of such technology resulted in a three primary benefits:

- Autonomous operations enabled by eliminating the need for hundreds of manual inputs in the form of control directives. Currently the task of creating the communications link is a manual and time-consuming process which requires operator input of approximately 700 control directives and the constant monitoring of several dozen displays to determine the exact execution status of the system.
- Reduced the level of expertise of an operator required to perform a communication track. Currently this complex process requires a high level of expertise from the operator, but through the development of the knowledge base by a domain expert this expertise is captured within the system itself.
- The knowledge base provides a declarative representation of operation procedures. Through the capture of this expertise the knowledge base documents the procedural steps of performing antenna communication services.

Related Work

There are a number of existing systems built to solve real-world planning or scheduling problems [15][16][17]. The problem of track plan generation combines elements from both these fields and thus traditional planners and schedulers cannot be directly applied. First, many classical planning elements must be addressed in this application such as subgoaling to achieve activity preconditions (e.g., the antenna must be "on_point" to lock up the receiver) and decomposing higher-level (abstract) activities into more detailed sub-activities. In addition, many scheduling elements are presents such as handling metric time and temporal constraints, and representing and reasoning about resources (e.g.,

receiver, antenna controller) and states (e.g., antenna position, subcarrier frequency, etc.) over time.

One other system has been designed to generate antenna track plans, the Deep Space Network Antenna Operations Planner (DPLAN) [4]. DPLAN utilizes a combination of AI hierarchical-task network (HTN) and operator-based planning techniques. Unlike DPLAN, ASPEN has a temporal reasoning system for expressing and maintaining temporal constraints and also has the capability for representing and reasoning about different types of resources and states. ASPEN can utilize different search algorithms such as constructive and repair-based algorithms, where DPLAN uses a best-first search. And, as described in the next section, ASPEN is currently being extended to perform dynamic planning for closed-loop error recovery, where DPLAN has only limited replanning capabilities.

As for the resource allocation type of scheduling performed by DANS, traditional scheduling system are not sufficient because of the unique type of constraints that the DSN scheduling problem poses.

A previous DSN scheduling system, OMP-26, was designed to perform the scheduling process for the smaller 9, 11, and 26 meter antennas. While the use of OMP-26 resulted in a five-fold reduction in scheduling labor and a doubling of network usage, OMP-26 does not have the ability to deal with the longer term forecasting required in the larger antenna network.

Another system developed for the resource allocation process in the DSN is the previously mentioned TIGRAS system. While TIGRAS has powerful tools for the visualization of network load and tools to assist with network forecasting, TIGRAS was not designed to perform automated rescheduling as was DANS nor demand access scheduling.

Future Work Providing Closed-Loop Control through Dynamic Planning

Currently, we are working on modifying and extending the current ASPEN Track Plan Generator to provide Closed Loop Execution and Recovery (CLEaR) for DSN track automation. CLEaR is built on top of CASPER [7], a real-time planning system built as an extension to ASPEN. The approach taken is to dynamically feed monitor data (sensor updates) back into the planning system as state updates. As these dynamic updates come in, the planning system verifies the validity of the current plan. If a violation is found in the plan, the system will perform local modification to construct a new valid plan.

Through this continual planning approach [6], the plan is disrupted as little as possible and the system is much more responsive and reactive to changes in the real (dynamic) world.

As part of the CLEaR effort further research is being done in the area *mixed-initiative control*. This addresses the interaction of an operator with, for all intensive purposes, an autonomous system. In these circumstances a planning and execution engine must maintain consistency with in the engine if an operator overrides the system so that once the operator returns the system to nominal operations the system is able to resume control without missing a heartbeat.

This CLEaR effort is also being integrated with a Fault Detection, Isolation and Recovery (FDIR) system. FDIR is an expert system providing monitor data analysis. As is often the case with large complex systems, monitor (sensor) data is often related in different ways that becomes difficult for a human to detect. The advantage of combining these two systems is that FDIR can first interpret the vast amount of data and summarize it into a set of meaningful values for a planning system to react to. We think of this union as intelligent analysis and intelligent response, much like a careful design and implementation; one without the other is of little use.

Conclusion

This paper has described the concept of the Deep Space Terminal (DS-T) and two of the key enable components in the DS-T autonomous operations capabilities. We first introduced the DSN problem domain and the DS-T 34-meter prototype antenna communications station. Next we described in detail the Demand Access Network Scheduling (DANS) system used to perform resource allocation/scheduling and the DS-T/ASPEN Script Generator used for antenna control script generation. We then concluded with results of the DS-T autonomous “lights out” operations demonstrations, discussion on related work, and presented some insight to future work being done in the area of DSN automation.

REFERENCES

- [1] Borden, C., Wang, Y-F., and Fox, G., “Planning and Scheduling User Services for NASA’s Deep Space Net-work,” *Working Notes of the 1997 International Workshop on Planning and Scheduling for Space Exploration and Science*, Oxnard, CA, 1997.
- [2] Chien, S., Decoste, D., Doyle, R., and Stolorz, P., “Making an Impact: Artificial Intelligence at the Jet Propulsion Laboratory,” *AI Magazine*, 18(1), 103-122, 1997.
- [3] (Chien et al. 1997b) Chien, S., Lam, R., and Vu, Q., “Resource Scheduling for a Network of Communications Antennas,” *Proceedings of the 1997 IEEE Aerospace Conference*, Aspen, CO, February, 1997.
- [4] Chien, S., Hill Jr., R., Govindjee, A., Wang, X., Estlin, T., A. Griesel, R. Lam and K. Fayyad, “A Hierarchical Architecture for Resource Allocation, Plan Execution, and Revision for Operation of a Network of Communication Antennas,” *Proceedings of the 1997 IEEE Conference on Robotics and Automation*, Albuquerque, NM, April 1997.
- [5] Chien, S., Fisher, F., Mortensen, H., Lo, E., Greeley, R., Govindjee, A., Estlin, T., and Wang, X., “Using Artificial Intelligence Planning Techniques to Automatically Reconfigure Software Modules,” in *Tasks and Methods in Applied Artificial Intelligence, Lecture Notes in Artificial Intelligence*, vol. 1416, Angel Pasqual del Pobil, Jose Mira, and Moonis Ali (Eds.), Springer-Verlag, 1998, pp. 415-426.
- [7] Chien, S., Knight, R., Stechert, A., Sherwood, R., and Rabideau, G., “Using Iterative Repair to Improve Responsiveness of Planning and Scheduling,” *IJCAI99 Workshop on Scheduling and Planning meet Real-time Monitoring in a Dynamic and Uncertain World*, Stockholm, Sweden, August 1999.
- [6] Chien, S., Knight, R., Stechert, A., Sherwood, R., and Rabideau, G., “Integrated Planning and Execution for Autonomous Spacecraft,” *Proceedings of the 1999 IEEE Aerospace Conference*, Aspen, CO, March 1999.
- [8] Deep Space Network, *Jet Propulsion Laboratory Publication* 400-517, April 1994.
- [9] Estlin, T., Fisher, F., Mutz, D., Chien, S., “Automated Planning for a Deep Space Communications Station”, *Proceedings of the IEEE Conference on Robotics and Automation*, Detroit, Michigan, May 1999.
- [10] Fisher, F., Chien, S., L. Paal, E. Law, N. Golshan, and M. Stockett, “An Automated Deep Space Communications Station,” *Proceedings of the IEEE Aerospace Conference*, Aspen, CO, March 1998.
- [11] Fisher, F., Mutz, D., Estlin, T., L. Paal, and Chien, S., “The Past, Present and Future of Ground Station Automation with in the DSN,” *Proceedings of the IEEE Aerospace Conference*, Aspen, CO, March 1999.
- [12] Fisher, F., Estlin, T., Mutz, D., and Chien, S., “Using Artificial Intelligence Planning to Generate Antenna Tracking Plans,” *Eleventh Conference on Innovative Applications of Artificial Intelligence*, Orlando, Florida, July 1999.
- [13] Fukunaga, A., Rabideau, G., Chien, S., and Yan, D., “Toward an Application Framework for Automated Planning and Scheduling,” *Proceedings of the 1997 International Symposium of Artificial Intelligence, Robotics and Automation for Space (ISAIRAS-97)*, Tokyo, Japan, July 1997
- [14] Hill, Jr., R., Chien, S., Fayyad, K., Smyth, C., Santos, T., and Bevan, R., “Sequence of Events Driven Automation of the Deep Space Network,” *Telecommunications and Data Acquisition* 42-124, October-December 1995.

[15] Tate, A., Drabble, B., and Kirby, R., "O-Plan2: An Open Architecture for Command Planning and Control," *Intelligent Scheduling* (Eds. M. Fox & M. Zweben), Morgan Kaufmann, 1994.

[16] Wilkins, D., "Extending the Classical AI Planning Paradigm," *Practical Planning*, Morgan Kaufmann, 1994.

[17] Zweben, M., Daun, B., Davis, E., and Deale, M., "Scheduling and Rescheduling with Iterative Repair," *Intelligent Scheduling*, Zweben, M., and Fox, M., eds., Morgan Kaufmann, 1994, pp.241-256.